

CLAIMS

What is claimed is:

1 1. A system for switching variable size packets in a network comprising:
2 at least one ingress controller which receives a plurality of packets and which segments each of
3 the packets into fixed sized fragments, the at least one ingress controller having a time-clock,
4 all ingress controller's time-clocks are synchronized to within a tolerance, each fragment is
5 tagged with at least a unique source of ID, a time-stamp, and a fragment-number to form a cell,
6 each cell belonging to one packet having the same time-stamp value, the ingress controller
7 sends each of the cells through a link such that a cell's destination is reachable through that
8 link;

9 a fabric element which receives cells from a plurality of inputs links, where such cells
10 are ordered, and sends ordered cells through a plurality of outputs, and through which the
11 destination of the cells is reachable, where the cell order is defined such that a cell ahead of
12 another either has a lagging time stamp, or if the timestamp is the same the cell ahead of
13 another has a source-id which has a predetermined priority, or if both the timestamp and the
14 source-id are the same the cell ahead of another has a lagging fragment-number; and

15 at least one egress controller which receives ordered cells from the plurality of input
16 links, and sends ordered cells through an output, where such order results in complete packets.

1 2. The system of claim 1 where the fabric element selects a cell for transmission
2 if all incoming links FIFOs have at least one cell.

1 3. The system of claim 1 where the output controller selects a cell for output if
2 all incoming links FIFOs have at least one cell.

1 4. The system of claim 1 where the ingress controller sends an empty cell
2 through output link if no data fragment is available, where such empty cells contain at least
3 the current local timestamp, and source-id.

1 5. The system of claim 2 where the ingress controller sends an empty cell
2 through output link if no data fragment is available, where such empty cells contain at least
3 the current local timestamp, and source-id.

1 6. The system of claim 4, where when empty cells are received by a fabric
2 element they are stored in the fabric element's corresponding link input FIFO if that FIFO
3 has less cells than a threshold.

1 7. The system of claim 5, where when empty cells are received by a fabric
2 element they are stored in the fabric element's corresponding link input FIFO if that FIFO
3 has less cells than a threshold.

1 8. The system of claim 1 where the fabric element sends an empty cell through
2 output link if no data fragment is available, where such empty cells contain at least the
3 timestamp, source-id, and fragment-number of the last cell that was processed by the sorter.

1
1 9. The system of claim 3 where the fabric element sends an empty cell through
2 output link if no data fragment is available, where such empty cells contain at least the
3 timestamp, source-id, and fragment-number of the last cell that was processed by the sorter.

1
2 10. The system of claim 6, where when an empty cell is received by the output
3 processor it is stored in the fabric element's corresponding link input FIFO if that FIFO has
less cells then a threshold.

1
2 11. The system of claim 7, where when an empty cell is received by the output
3 processor it is stored in the fabric element's corresponding link input FIFO if that FIFO has
less cells then a threshold.

1
2 12. The system of claim 1, where the fabric element selects a cell for
3 transmission if all active incoming links FIFOs have at least one cell, where active links are
links through which a cell with at least a timestamp, source-id was received in a past period.

1
2 13. The system of claim 1, where the output controller selects a cell for output if
3 all active incoming links FIFOs have at least one cell, where active links are links through
which a cell with at least a timestamp, source-id was received in a past period.

1
2 14. The system of claim 1, where ingress controller sends cells through a link
such that cell's destination is reachable through that link and such that all possible links are

3 load balanced.

1 15. The system of claim 1, where fabric element sends cells through links
2 through which a cell's destination is reachable, and where all possible output links for such
3 cell are load balanced.

1 16. The system of claim 14, where the ingress controller maintains a reachability
2 table used to determine which links are possible for a destination.

17. The system of claim 15, where the fabric element maintains a reachability
table used to determine which links are possible for a destination.